arXiv:1505.06356v1 [stat.CO] 23 May 2015

Particle ancestor sampling for near-degenerate or intractable state transition models*

Fredrik Lindsten, Pete Bunch, Sumeetpal S. Singh,
and Thomas B. Schön

23 May 2015

Abstract

We consider Bayesian inference in sequential latent variable models in general, and in nonlinear state space models in particular (i.e., state smoothing). We work with sequential Monte Carlo (SMC) algorithms, which provide a powerful inference framework for addressing this problem. However, for certain challenging and common model classes the state-of-the-art algorithms still struggle. The work is motivated in particular by two such model classes: *(i)* models where the state transition kernel is (nearly) degenerate, i.e. (nearly) concentrated on a low-dimensional manifold, and *(ii)* models where point-wise evaluation of the state transition density is intractable. Both types of models arise in many applications of interest, including tracking, epidemiology, and econometrics. The difficulties with these types of models is that they essentially rule out forward-backward-based methods, which are known to be of great practical importance, not least to construct computationally efficient particle Markov chain Monte Carlo (PMCMC) algorithms. To alleviate this, we propose a "particle rejuvenation" technique to enable the use of the forward-backward strategy for (nearly) degenerate models and, by extension, for intractable models. We derive the proposed method specifically within the context of PMCMC, but we emphasise that it is applicable to any forward-backward-based Monte Carlo method.

1 Problem formulation

State space models (SSMs) are widely used for modelling time series and dynamical systems. A general, discrete-time SSM can be written as

$$x_t \,|\, x_{t-1} \sim f(x_t \,|\, x_{t-1}), \tag{1a}$$

$$y_t \,|\, x_t \sim g(y_t \,|\, x_t), \tag{1b}$$

*Supported by the projects *Learning of complex dynamical systems* (Contract number: 637-2014-466) and *Probabilistic modeling of dynamical systems* (Contract number: 621-2013-5524), both funded by the Swedish Research Council, and the project *Bayesian Tracking and Reasoning over Time* (Reference: EP/K020153/1), funded by the EPSRC.

where $x_t \in \mathsf{X}$ is the latent state, $y_t \in \mathsf{Y}$ is the observation (both at time t) and $f(\cdot)$ and $g(\cdot)$ are probability density functions (PDFs) encoding the state transition and the observation likelihood, respectively. The initial state is distributed according to $x_1 \sim \mu(x_1)$.

Statistical inference in SSMs typically involves computation of the *smoothing distribution*, that is, the posterior distribution of a sequence of state variables $X_T := (x_1, \ldots, x_T) \in \mathsf{X}^T$ conditionally on a sequence of observations $Y_T := (y_1, \ldots, y_T) \in \mathsf{Y}^T$. The smoothing distribution plays a key role both for offline (batch) state inference and for system identification via data augmentation methods, such as expectation maximisation [15] and Gibbs sampling [43].

The motivation for the present work comes from two particularly challenging classes of SSMs:

(M1) If the state transition kernel $f(\cdot)$ of the system puts all probability mass on some low-dimensional manifold, we say that the transition is *degenerate* (for simplicity we use "probability density notation" even in the degenerate case). Degenerate transition kernels arise, e.g., if the dynamical evolution is modeled using additive process noise with a rank-deficient covariance matrix. Models of this type are common in certain application areas, e.g., navigation and tracking; see [26] for several examples. Likewise, if $f(\cdot)$ is concentrated around a low-dimensional manifold (i.e., the transition is highly informative, or the process noise is small) we say that the transition is *nearly degenerate*.

(M2) If the state transition density function $f(\cdot)$ is not available on closed form, the transition is said to be *intractable*. The typical scenario is that $f(\cdot)$ is a regular (non-degenerate) PDF which it is possible to simulate from, but which nevertheless is intractable. At first, this scenario might seem contrived, but it is in fact quite common in practice. In particular, whenever the dynamical function is defined implicitly by some computer program or black-box simulator, or as a "complicated" nonlinear transformation of a known noise input, it is typically not possible to explicitly write down the corresponding transition PDF; see, e.g., [36, 3] for examples.

The main difficulty in performing inference for model classes **(M1)** and **(M2)** lies in that the so-called *backward kernel* of the model (see, e.g., [30]),

$$p(x_t \,|\, x_{t+1}, y_{1:t}) \propto f(x_{t+1} \,|\, x_t) p(x_t \,|\, y_{1:t}), \tag{2}$$

will also be (nearly) degenerate or intractable, respectively. This is problematic since many state-of-the-art methods rely on the backward kernel for inference; see, e.g., [24, 31, 47]. In particular, the backward kernel is used to implement the well-known forward-backward smoothing strategy. We will come back to this in the subsequent sections when we discuss the details of these inference methods and how we propose an extension to the methodology geared toward this issue.

The main contribution of this paper is constituted by a construction allowing us to replace the simulation from the (problematic) backward kernel, with

simulation from a joint density over a subset of the future state variables. Importantly, this joint density will typically have more favourable properties than the backward kernel itself, whenever the latter is (nearly) degenerate. Simulating from the joint density results in a *rejuvenation* of the state values, which intuitively can be understood as a bridging between past and future state variables. Furthermore, to extend the scope of this technique we propose a nearly degenerate approximation of an intractable transition model. Using this approximation we can thus use the proposed particle rejuvenation strategy to perform inference also in models with intractable transitions.

2 Background, methodology, and related work

2.1 Computational inference

The strong assumptions of linearity and Gaussianity that were originally invoked for state space inference have been significantly weakened by the development of computational statistical methods. Among these, Markov chain Monte Carlo (MCMC) and sequential Monte Carlo (SMC, a.k.a. particle filter) methods play prominent roles. MCMC methods (see e.g., [41, 2]) are based on simulating a Markov chain which admits the target distribution of interest—here, the joint smoothing distribution—as its unique stationary distribution. SMC methods, on the other hand (see, e.g., [16, 13, 12]), use a combination of sequential importance sampling [27] and resampling [42] to approximate a sequence of probability distributions defined on a sequence of measurable spaces of increasing dimension. This is accomplished by approximating each target distribution by an empirical point-mass distribution based on a collection of random samples, referred to as *particles*, with corresponding non-negative *importance weights*. For instance, the target distributions of an SMC sampler can comprise the smoothing distributions for an SSM for $t = 1, \ldots, T$ and, indeed, SMC methods have emerged as a key tool for approximating the flow of smoothing distributions for general SSMs.

In addition to these methods, we have in recent years seen much interest in the combination of SMC and MCMC in so-called particle MCMC (PMCMC) methods [3]. These methods are based on an auxiliary variable construction which opens up the use of SMC (or variants thereof) to construct efficient, high-dimensional MCMC transition kernels. The introduction of PMCMC has spurred intensive research in the community spanning methodological [9, 31, 48, 38], theoretical [32, 4, 14, 10, 17, 1], and applied [46, 39, 25, 40] work.

In this paper we consider in particular the *particle Gibbs* (PG) algorithm, introduced by Andrieu et al. [3]. The PG algorithm relies on running a modified particle filter, in which one particle trajectory is set deterministically according to an *a priori* specified reference trajectory (see Section 3). After a complete run of the SMC algorithm, a new trajectory is obtained by selecting one of the particle trajectories with probabilities given by their importance weights. This results in a Markov kernel on the *space of trajectories* X^T. Interestingly, the

conditioning on a reference trajectory ensures that the limiting distribution of this Markov kernel is exactly the target distribution of the sampler (e.g., the joint smoothing distribution) for any number of particles $N \geq 2$ used in the underlying SMC algorithm. The PG algorithm can be interpreted as a Gibbs sampler for the extended model where the random variables generated by the SMC sampler are treated as auxiliary variables.

Like any Gibbs sampler, PG has the advantage over Metropolis-Hastings of not requiring an accept/reject stage. However, the resulting chain is still liable to mix (very) slowly if the particle filter suffers from path-space degeneracy [10, 30]. Unfortunately, path degeneracy is inevitable for high-dimensional (large T) problems, which significantly reduces the applicability of PG. This problem has been addressed in a generic setting by adding additional sampling steps to the PG sampler, either during the filtering stage, known as particle Gibbs with ancestor sampling (PGAS) [31], or in an additional backward sweep, known as particle Gibbs with backward simulation (PGBS) [48, 47]. The improvement arises from sampling new values for individual particle ancestor indexes, and thus allowing the reference trajectory to be updated gradually which can mitigate the effects of path degeneracy. It has been found that this can *vastly* improve mixing, making the method much more robust to a small number of particles as well as growth in the size of the data.

2.2 Tackling (nearly) degenerate or intractable transitions

A problem with PGAS and PGBS, however, is that they rely on a particle approximation of the *backward kernel* for updating the ancestry of the particles. For an SSM the backward kernel is given by (2), which implies that if $f(\cdot)$ is (nearly) degenerate or intractable, then so is the backward kernel. As an effect, the probability of sampling any change in the particle ancestry is low. This largely removes the effect of the extra sampling steps introduced in order to reduce path degeneracy. In fact, if the transition is truly degenerate, then the probability of updating the ancestry will be exactly zero. Intuitively, the problem is that the only state history consistent with a particular "future state" is that from which the future state was originally generated.

To mitigate this effect we propose to use a procedure which we call *particle rejuvenation*. A specific instance of this technique has previously been used together with forward/backward particle smoothers [6]—when sampling an ancestor index they simultaneously sample a new value for the associated state. In the preliminary work [7] we investigated the effect of this approach on the PGBS sampler. This opened up for steering the potential state histories towards the fixed future, consequently increasing the probability of changing the ancestry and thus improving the mixing of the Markov chain. Independently, Carter et al. [9] have derived essentially the same method. However, they introduce a different extended target distribution in order to justify this addition, which necessitates changes to the underlying SMC sampler, whereas our developments allows us to use the standard PMCMC construction.

In the present work we extend the method from [7, 9] by considering a more

generic setting, in which we propose to rejuvenate any "subset" of the reference trajectory along with the particle ancestry (Section 4). This provides additional flexibility over previous approaches. As we illustrate in Section 5, this increased flexibility is necessary in order to address the challenges associated with several models of interest containing (nearly) degenerate transitions (type **(M1)**). Furthermore, in Section 6 we show how an intractable transition can be approximated with a nearly degenerate one. Employing the particle rejuvenation strategy to this approximation results in inference methods applicable to models of type **(M2)** where $f(\cdot)$, and thus the backward kernel (2), is intractable. Here, again, the increased flexibility provided by the generalised particle rejuvenation strategy is key. We work specifically with the PGAS algorithm, but all the proposed modifications are also applicable to PGBS and, in fact, to any other backward-simulation-based method (see [30]).

3 Particle Gibbs with Ancestor Sampling

While SSMs comprise one of the main motivating factors for the development of SMC and PMCMC (and we shall return to these models in the sequel), these methods can be applied more generally, see, e.g., [3, 30]. For this reason we will carry out the derivation in a more general setting.

We start by reviewing the PGAS algorithm by Lindsten et al. [31]. Consider the problem of sampling from a possibly high-dimensional probability distribution. As before, we write $X_T := (x_1, \ldots, x_T) \in \mathsf{X}^T$ for the random variables of the model. Let $\gamma_t(X_t)$ for $t = 1, \ldots, T$ be a sequence of unnormalised densities on X^t, which we assume can be evaluated point-wise. Let $\bar{\gamma}_t(X_t)$ be the corresponding normalised PDFs. The central object of interest is assumed to be the final PDF in the sequence: $\bar{\gamma}_T(X_T)$. For an SSM we would typically have $\bar{\gamma}_t(X_t) = p(X_t \,|\, Y_t)$ and $\gamma_t(X_t) = p(X_t, Y_t)$.

The PGAS algorithm [31] is a procedure for constructing a Markov kernel on X^T which admits $\bar{\gamma}_T(X_T)$ as its unique stationary distribution. Consequently, by simulating a Markov chain with transition probability given by the, so-called, PGAS kernel, we will (after an initial transient phase) obtain samples distributed according to $\bar{\gamma}_T(X_T)$. The PGAS kernel can thus readily be used in MCMC and related Monte Carlo methods based on Markov kernels. As mentioned above, PGAS is an instance of PMCMC. Specifically, the construction of the PGAS kernel relies on (a variant of) an SMC sampler, targeting the sequence of intermediate distributions $\gamma_t(X_t)$ for $t = 1, \ldots, T$.

For later reference we introduce some additional notation. First, we will make frequent use of the notation already exemplified above, where capital letters are used for sequences of variables, e.g., by writing X_t for the "past" history of the state sequence. Similarly, $\widetilde{X}_{t+1} := X_T \setminus X_t = (x_{t+1}, \ldots, x_T)$ denotes the "future" state sequence. Here, we have used *set notation* to refer to a collection of latent variables and we will make frequent use also of this notation in the sequel. In particular, we will write $\Xi \subseteq X_T$ for an arbitrary subset of the latent random variables of the model, which could refer to individual

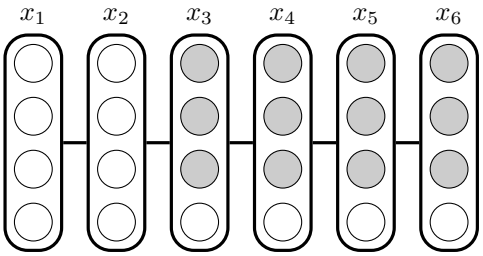

Figure 1: Illustration of subset notation $\Xi \subset X_T$. Here, $\mathsf{X} = \mathbb{R}^4$ and each component of the vector x_t is illustrated by a circle. The gray disks illustrate a subset $\Xi \subset X_6$, consisting of $\Xi = \{x_{j,t} : j = 1, 2, 3 \text{ and } t = 3, 4, 5, 6\}$.

components of x_t, say, when x_t is a vector (see Section 6). For clarity, we exemplify this usage of set notation in Figure 1.

The PGAS algorithm is reminiscent of a standard SMC sampler, but with the important difference that one particle at each iteration is specified *a priori*. These particles, denoted as $(x'_1, \ldots, x'_T) =: X'_T$ and with corresponding particle indexes $(b_1, \ldots, b_T) =: B_T$, serve as a reference for the sampler, as detailed below.

As in a standard SMC sampler, we approximate the sequence of target densities $\bar{\gamma}_t(X_t)$ for $t = 1, \ldots, T$ by collections of weighted particles. Let $\{X^i_{t-1}, w^i_{t-1}\}_{i=1}^N$ be a particle system approximating $\bar{\gamma}_{t-1}(X_{t-1})$ by the empirical distribution,

$$\bar{\gamma}^N_{t-1}(X_{t-1}) := \sum_{i=1}^N \frac{w^i_{t-1}}{\sum_{j=1}^N w^j_{t-1}} \delta_{X^i_{t-1}}(X_{t-1}). \tag{3}$$

Here, $\delta_z(x)$ is a point mass located at z. To propagate this particle system to time t, we introduce the auxiliary variables $\{a^i_t\}_{i=1}^N$, referred to as *ancestor indexes*, encoding the genealogy of the particle system. More precisely, x^i_t is generated by first sampling its ancestor,

$$\mathbb{P}(a^i_t = j) \propto w^j_{t-1}, \qquad\qquad j = 1, \ldots, N. \tag{4}$$

Then, x^i_t is drawn from some proposal distribution,

$$x^i_t \sim r_t(\cdot \mid X^{a^i_t}_{t-1}). \tag{5}$$

When we write X^i_t we refer to the ancestral path of particle x^i_t. That is, the particle trajectory is defined recursively as

$$X^i_t := (X^{a^i_t}_{t-1}, x^i_t). \tag{6}$$

In this formulation the resampling step is implicit and corresponds to sampling the ancestor indexes. Finally, the particle is assigned a new importance weight:

$w_t^i = \omega_t(X_t^i)$ where the weight function is given by

$$\omega_t(X_t) = \frac{\gamma_t(X_t)}{\gamma_{t-1}(X_{t-1}) r_t(x_t \mid X_{t-1})}. \tag{7}$$

In a standard SMC sampler, the procedure above is repeated for each $i = 1, \ldots, N$, to generate N particles at time t. However, as mentioned above, the PGAS procedure relies on keeping one particle fixed at each iteration in order to obtain the correct limiting distribution of the resulting MCMC kernel. This is accomplished by simulating according to (4) and (5) only for $i \in \{1, \ldots, N\} \setminus b_t$. The final particle is set deterministically according to the reference: $x_t^{b_t} = x'_t$.

To be able to construct the particle trajectory $X_t^{b_t}$ as in (6), the reference particle has to be associated with an ancestor at time $t-1$. This is done by *ancestor sampling*; that is, we simulate randomly a value for the corresponding ancestor index $a_t^{b_t}$. Lindsten et al. [31] derive the ancestor sampling distribution, and show that $a_t^{b_t}$ should be simulated according to,

$$\mathbb{P}(a_t^{b_t} = i) \propto w_{t-1}^i \frac{\gamma_T(X_{t-1}^i \cup \widetilde{X}'_t)}{\gamma_{t-1}(X_{t-1}^i)}, \qquad i = 1, \ldots, N. \tag{8}$$

Note that the expression depends on the complete "future" reference path $\widetilde{X}'_t = (x'_t, \ldots, x'_T)$. In the above, $X_{t-1}^i \cup \widetilde{X}'_t$ refers to the complete path formed by concatenating the two partial trajectories.

The expression above can be understood as an application of Bayes' theorem; the weight w_{t-1}^i is the prior probability of particle X_{t-1}^i and the density ratio corresponds to the likelihood of "observing" $\widetilde{X}'_t$ given X_{t-1}^i. The actual motivation for the ancestor sampling distribution (8), however, relies on a collapsing argument [31]. We will review this idea in the subsequent section where we propose a generalisation of the technique to increase the flexibility of the PGAS algorithm.

Finally, after a complete pass of the modified SMC procedure outlined above for $t = 1, \ldots, T$, a new trajectory $X_T^\star$ is sampled by selecting among the particle trajectories with probability given by their corresponding importance weights. That is, we sample k with $\mathbb{P}(k = i) \propto w_T^i$, $i = 1, \ldots, N$ and return $X_T^\star = X_T^k$. We summarise the PGAS procedure in Algorithm 1. Note that the algorithm stochastically simulates the trajectory $X_T^\star$ conditionally on the reference trajectory X'_T, thus implicitly defining a Markov kernel on X^T.

Remark 1. As pointed out in [10, 31], the indexes B_T are nuisance variables that are unnecessary from a practical point of view, i.e., it is possible to set B_T to any arbitrary convenient sequence, e.g. $B_T = (N, \ldots, N)$. We use this convention in Algorithms 1 and 2. However, explicit reference to the index variables will simplify the derivation of the algorithm.

Remark 2. PGAS is a variation of the PG sampler by Andrieu et al. [3]. Algorithmically, the only difference between the methods lies in the ancestor sampling step (8), where, in the original PG sampler we would simply set $a_t^N = N$

Algorithm 1 PGAS Markov kernel [31]

Require: Reference trajectory $X'_T \in \mathsf{X}^T$.

1: Set $x_1^N = x'_1$.
2: Draw $x_1^i \sim r_1(\cdot)$ for $i = 1, \ldots, N-1$.
3: Set $w_1^i = \gamma_1(x_1^i)/r_1(x_1^i)$ for $i = 1, \ldots, N$.
4: **for** $t = 2$ **to** T **do**
5: Simulate (a_t^i, x_t^i) as in (4, 5) for $i = 1, \ldots, N-1$.
6: Simulate a_t^N according to (8).
7: Set $x_t^N = x'_t$.
8: Set $X_t^i = (X_{t-1}^{a_t^i}, x_t^i)$ for $i = 1, \ldots, N$.
9: Set $w_t^i = \omega_t(X_t^i)$ for $i = 1, \ldots, N$.
10: **end for**
11: Draw k with $\mathbb{P}(k = i) \propto w_T^i$.
12: **return** $X_T^\star = X_T^k$.

(deterministically). While this is a small modification, it can have a very large impact on the convergence speed of the method, as discussed and illustrated in [31]. Informally, the reason for this is that the *path degeneracy* of SMC samplers will cause the PG sampler to degenerate toward the reference trajectory. That is, with high probability $x_s^\star = x'_s$ for any $s \ll T$, effectively causing the sampler to be stuck at its initial value in certain parts of the state space X^T. Ancestor sampling mitigates this issue by assigning a new ancestry of the reference trajectory via (8). Path degeneracy will still occur, but the sampler tends to degenerate toward something different than the reference trajectory ($x_s^\star \neq x'_s$ with non-negligible probability) enabling much more efficient exploration of the state space.

4 Particle rejuvenation for the PGAS kernel

We now turn to the new procedure: a *particle rejuvenation* strategy for PGAS. The idea is to update the ancestor index $a_t^{b_t}$ jointly with some subset $\Xi_t \subseteq \widetilde{X}'_t$ of the reference trajectory. The intuitive motivation is that by "loosening up" the reference trajectory, we get more freedom of changing its ancestry. Before we continue with the derivation of the particle rejuvenation strategy, however, we shall consider a simple numerical example to motivate the development.

4.1 An illustrative example

Suppose we wish to track the motion of an object in 3D space from noisy measurements of its position. For the transition model $f_\theta(x_t \,|\, x_{t-1})$ we assume near constant velocity motion [29], and the observation model is noisy measurements of bearing, elevation and range. Here we have also introduced an unknown system parameter θ corresponding to the scale factor on the transition covariance

matrix, which characterises the target manoeuvrability. Based on a batch of $T = 100$ observations Y_T, we wish to learn the unknown parameter θ by Gibbs sampling, iteratively simulating from $p(\theta \mid X_T, Y_T)$ and $p(X_T \mid \theta, Y_T)$. (See [7] for more details on the model and experimental setup.)

This model can be problematic for PGAS if the scale parameter θ is small, since this implies a highly informative, i.e., nearly degenerate, transition model. To be more precise, for an SSM, with the unnormalised target density being $\gamma_{\theta,t}(X_t) = p(X_t, Y_t \mid \theta)$, it follows that the ancestor sampling distribution in (8) simplifies to

$$\mathbb{P}(a_t^{b_t} = i) = \frac{w_{t-1}^i f_\theta(x_t' \mid x_{t-1}^i)}{\sum_{j=1}^N w_{t-1}^j f_\theta(x_t' \mid x_{t-1}^j)}, \qquad i = 1, \ldots, N. \tag{9}$$

Now, if θ—the process noise—is small, then $f_\theta(x_t \mid x_{t-1})$ is largely concentrated on a single point. Hence, the distribution in (9) will also be concentrated on one value and there is little freedom in changing the ancestry of x_t' at time $t-1$. The result is that the effect of ancestor sampling is diminished.

The idea with particle rejuvenation is that simultaneously sampling a new state x_t', for instance, jointly with the ancestor index opens up for *bridging* between the states x_{t-1}^i and x_{t+1}'. This leads to a substantially higher probability of updating the ancestor indexes during each iteration, and hence faster mixing.

Using this rejuvenation strategy roughly doubles the computation required to execute each iteration of PGAS. In simulations, the median factor of improvement in the probability of accepting an ancestor change at each time step is 2.4 (inter-quartile range 1.9–2.8). By contrast, simply doubling the number of particles results in a median factor of improvement of only 1.5 (inter-quartile range 1.4–1.6). The difference is even more pronounced in terms of autocorrelation, as shown in Figure 2. (We also ran the basic PG algorithm, without ancestor sampling, but due to path degeneracy the method did not converge and the results are therefore not reported here.)

4.2 Extended target distribution

The formal motivation for the validity of PMCMC algorithms is based on an auxiliary variables argument. More precisely, Andrieu et al. [3] introduce an *extended target distribution* which is defined on the space of all the random variables generated by the run of an SMC algorithm. Let

$$\mathbf{x}_t := \{x_t^1, \ldots, x_t^N\}, \qquad \text{and} \qquad \mathbf{a}_t := \{a_t^1, \ldots, a_t^N\},$$

denote the particles and ancestor (resampling) indexes generated at time t, respectively. We also write $\mathbf{X}_t := \{\mathbf{x}_1, \ldots, \mathbf{x}_t\}$ and $\mathbf{A}_t := \{\mathbf{a}_2, \ldots, \mathbf{a}_t\}$. Furthermore, let k be the index of one specific reference trajectory. To make the particle indexes of the reference trajectory X_T^k explicit we define recursively: $b_T = k$ and $b_t = a_{t+1}^{b_{t+1}}$ for $t < T$. Hence, b_t corresponds to the index of the reference particle at time t, obtained by tracing the ancestry of x_T^k. If follows that $X_T^k = (x_1^{b_1}, \ldots, x_T^{b_T})$.

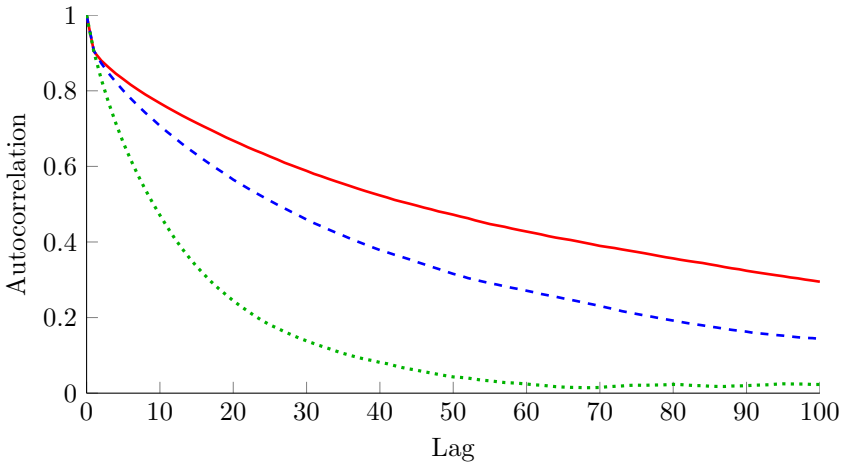

Figure 2: Autocorrelation functions for the parameter θ (scale factor of transition covariance). Averages over 5 runs of PGAS without rejuvenation using $N = 100$ particles (solid red) and $N = 200$ particles (dashed blue), and with rejuvenation using $N = 100$ particles (dotted green).

The extended target distribution for PMCMC samplers is then given by (see [3])

$$\pi_T^N(\mathbf{X}_T, \mathbf{A}_T, k) := \frac{\bar{\gamma}_T(X_T^k)}{N^T} \left\{ \prod_{\substack{i=1 \\ i \neq b_1}}^{N} r_1(x_1^i) \right\} \prod_{t=2}^{T} \left\{ \prod_{\substack{i=1 \\ i \neq b_t}}^{N} \frac{w_{t-1}^{a_t^i}}{\sum_{j=1}^{N} w_{t-1}^j} r_t(x_t^i \mid X_{t-1}^{a_t^i}) \right\}. \tag{10}$$

A key property of this distribution is that it admits the original target distribution $\bar{\gamma}_T$ as a marginal. That is, if $(\mathbf{X}_T, \mathbf{A}_T, k)$ are distributed according to π_T^N, then the marginal distribution of X_T^k is $\bar{\gamma}_T$. This implies that π_T^N can be used in place of $\bar{\gamma}_T$ in an MCMC scheme; this is the technique used by PMCMC samplers.

4.3 Partial collapsing and particle rejuvenation

In particular, the PGAS algorithm that we reviewed in Section 3 corresponds to a partially collapsed Gibbs sampler for the extended target distribution π_T^N. The complete Gibbs sweep corresponding to Algorithm 1 is given in the appendix. Here, however, we will focus on the ancestor sampling step (8). As mentioned in Remark 2, this step is very useful for improving the mixing of the PG algorithm. However, as described above, for certain classes of models the likelihood of the "future" reference path $\widetilde{X}_t'$ can be very low under alternative histories $\{X_{t-1}^i\}_{i=1}^N$. The PDF ratio in expression (8) will thus cause the ancestor sampling distribution to be highly concentrated on $i = b_{t-1}$, i.e., $\mathbb{P}(a_t^{b_t} = b_{t-1}) \approx 1$.

In particular, this is true for nearly degenerate SSMs; recall model class **(M1)**. In fact, for truly degenerate models it may be that $\mathbb{P}(a_t^{b_t} = b_{t-1}) = 1$, which implies that the ancestor sampling step has no effect and PGAS is reduced to the basic PG scheme.

Observe that, while the PGAS algorithm attempts to update the *ancestry* of the reference particles X'_T, it does not update the values of the particles themselves. This observation can be used to mitigate the aforementioned shortcoming of the algorithm, as we will now illustrate. The proposed modification is conceptually simple, but its practical implications for improving the mixing of the PGAS algorithm can be quite substantial for many models of interest.

The idea is to *simultaneously* update the ancestor index $a_t^{b_t}$ together with a part of the future reference trajectory $\widetilde{X}'_t$. This results in an increased flexibility of bridging the future reference path with an alternative history, thereby increasing the probability of changing its ancestry. By "a part of", we here refere to any collection of random variables $\Xi_t \subseteq \widetilde{X}_t$ (see Figure 1 for an illustration). Typically, the larger this subset is, the larger will the increased flexibility be. However, this has to be traded off with the difficulty of updating Ξ_t, which can be substantial if Ξ_t is overly high-dimensional.

For notational convenience, let $\widetilde{\Xi}_t = \widetilde{X}_t \setminus \Xi_t$. The AS step of the PGAS algorithm is then replaced by a step where we simulate $(a_t^{b_t}, \Xi_t)$ jointly from the conditional distribution on $\{1, \ldots, N\} \times \text{range}(\Xi_t)$:

$$\pi_T^N(a_t^{b_t}, \Xi_t \,|\, \mathbf{X}_{t-1}, \mathbf{A}_{t-1}, \widetilde{\Xi}_t, \widetilde{B}_t), \tag{11}$$

where $\widetilde{B}_t := (b_t, \ldots, b_T)$. Note that this is a so-called *partially collapsed* Gibbs move, since not all the non-simulated variables of the model are conditioned upon. Specifically, we have excluded all the future particles and ancestor indexes, except for those corresponding to the reference path. The justification for this is that we are sampling conceptually from the distribution

$$\pi_T^N(a_t^{b_t}, \Xi_t, \{\widetilde{\mathbf{X}}_t \setminus \widetilde{X}_t\}, \{\widetilde{\mathbf{A}}_t \setminus \widetilde{B}_t\} \,|\, \mathbf{X}_{t-1}, \mathbf{A}_{t-1}, \widetilde{\Xi}_t, \widetilde{B}_t), \tag{12}$$

which is a standard Gibbs update for the extended target distribution (10). However, no consecutive operation will depend on the variables $\widetilde{\mathbf{X}}_t \setminus \widetilde{X}_t$ and $\widetilde{\mathbf{A}}_t \setminus \widetilde{B}_t$, which has the implication that these variables need not be generated at all; see [19].

Following [31] we obtain an expression for the conditional distribution (11) which much resembles the original ancestor sampling distribution (8), namely,

$$\pi_T^N(a_t^{b_t}, \Xi_t \,|\, \mathbf{X}_{t-1}, \mathbf{A}_{t-1}, \widetilde{\Xi}'_t, \widetilde{B}_t) \propto w_{t-1}^{a_t^{b_t}} \frac{\gamma_T(X_{t-1}^{a_t^{b_t}} \cup \Xi_t \cup \widetilde{\Xi}'_t)}{\gamma_{t-1}(X_{t-1}^{a_t^{b_t}})}. \tag{13}$$

Note, however, that this is a distribution on $\{1, \ldots, N\} \times \text{range}(\Xi_t)$, and simulating from this distribution allows us to update the ancestor index $a_t^{b_t}$ jointly with a part of the reference trajectory Ξ_t. Additionally, at time $t = 1$ we can

update Ξ_1 by simulating from the conditional distribution

$$\pi_T^N(\Xi_1 \,|\, \widetilde{\Xi}_1', \widetilde{B}_1) \propto \gamma_T(\Xi_1 \cup \widetilde{\Xi}_1'). \tag{14}$$

In most cases, exact simulation from (13) or (14) is not possible. However, this issue can be dealt with by instead simulating from some MCMC kernels leaving these distributions invariant, resulting in a standard combination of MCMC samplers, see e.g. [45]. Hence, let K_t denote a Markov kernel on $\{1, \ldots, N\} \times \text{range}(\Xi_t)$ which leaves the distribution (13) invariant (sampling from (14) at time $t = 1$ follows analogously). The proposed modified PGAS method is then given by Algorithm 2.

Algorithm 2 PGAS with particle rejuvenation

Require: Reference trajectory $X_T' \in \mathsf{X}^T$.

1: Simulate $\Xi_1^\star \sim K_1(\Xi_1', \cdot)$ and update X_T' accordingly: $X_T' \leftarrow \{X_T' \setminus \Xi_1'\} \cup \Xi_1^\star$.
2: Set $x_1^N = x_1'$.
3: Draw $x_1^i \sim r_1(\cdot)$ for $i = 1, \ldots, N-1$.
4: Set $w_1^i = \gamma_1(x_1^i)/r_1(x_1^i)$ for $i = 1, \ldots, N$.
5: **for** $t = 2$ **to** T **do**
6: Simulate (a_t^i, x_t^i) as in (4, 5) for $i = 1, \ldots, N-1$.
7: Simulate $(a_t^N, \Xi_t^\star) \sim K_t((N, \Xi_t'), \cdot)$ and update X_T' accordingly: $X_T' \leftarrow \{X_T' \setminus \Xi_t'\} \cup \Xi_t^\star$.
8: Set $x_t^N = x_t'$.
9: Set $X_t^i = (X_{t-1}^{a_t^i}, x_t^i)$ for $i = 1, \ldots, N$.
10: Set $w_t^i = \omega_t(X_t^i)$ for $i = 1, \ldots, N$.
11: **end for**
12: Draw k with $\mathbb{P}(k = i) \propto w_T^i$.
13: **return** $X_T^\star = X_T^k$.

Remark 3. The previous particle rejuvenation strategies proposed independently by us [7] and Carter et al. [9] correspond to the special case obtained by setting $\Xi_t = x_t$. However, as we shall see in Sections 5 and 6, this is insufficient in many cases. In particular, to address the challenges associated with some degenerate models **(M1)** and models with intractable transitions **(M2)**, we need additional flexibility in selecting Ξ_t. Furthermore, a difference between the current derivation and the one presented by Carter et al. [9] is that they do not make use of the technique of *partial collapsing*. As a consequence, they are forced to re-define the extended target distribution (10), resulting in (unnecessary) modifications of the SMC scheme and it implies that their approach is only applicable when using an explicit backward pass (as in PGBS).

Remark 4. While ergodicity of the kernels K_t is of practical importance in order to obtain a large performance improvement from the 'ancestor sampling & particle rejuvenation' strategy, it is not needed to guarantee ergodicity of the overall sampling scheme. In particular, if $K_t((a_t^{b_t}, \Xi_t), \cdot) = \delta_{(a_t^{b_t}, \Xi_t)}(\cdot)$, the proposed method reduces to the original PG algorithm by [3] which is known to be uniformly geometrically ergodic under weak assumptions [32, 4].

Below, we present two specific techniques for designing the kernels K_t that can be useful in the present context.

Metropolis-Hastings (MH) We can target (13) using MH. From current values (a'_t, Ξ'_t), we can propose new values $(a^\star_t, \Xi^\star_t)$ by drawing from,

$$\frac{\nu_{t-1}^{a_t}}{\sum_{j=1}^N \nu_{t-1}^j} \phi_t(\Xi_t \mid X_{t-1}^{a_t}, \Xi'_t, \widetilde{\Xi}'_t), \tag{15}$$

where $\{\nu_t^i\}_{i=1}^N$ is a set of proposal weights for the ancestor index and ϕ_t is a proposal density for rejuvenating the reference particles Ξ'_t. The resulting acceptance probability is then,

$$\begin{aligned}
&\alpha\left\{(a'_t, \Xi'_t) \to (a^\star_t, \Xi^\star_t)\right\} \\
&= \min\left\{1, \frac{\nu_{t-1}^{a'_t}\phi_t(\Xi'_t \mid X_{t-1}^{a'_t}, \Xi^\star_t, \widetilde{\Xi}'_t)}{\nu_{t-1}^{a^\star_t}\phi_t(\Xi^\star_t \mid X_{t-1}^{a^\star_t}, \Xi'_t, \widetilde{\Xi}'_t)} \frac{w_{t-1}^{a^\star_t}\gamma_T(X_{t-1}^{a^\star_t} \cup \Xi^\star_t \cup \widetilde{\Xi}'_t)}{w_{t-1}^{a'_t}\gamma_T(X_{t-1}^{a'_t} \cup \Xi'_t \cup \widetilde{\Xi}'_t)} \frac{\gamma_{t-1}(X_{t-1}^{a'_t})}{\gamma_{t-1}(X_{t-1}^{a^\star_t})}\right\}.
\end{aligned} \tag{16}$$

Conditional importance sampling Given that we are working within the PG framework, a more natural approach might be to use a conditional importance sampling (CIS) Markov kernel. This can be viewed simply as an instance of the PG kernel applied to a single time step. Consider an importance sampling proposal distribution for (a_t, Ξ_t) (cf. (15)),

$$\frac{\nu_{t-1}^{a_t}}{\sum_{j=1}^N \nu_{t-1}^j} \psi_t(\Xi_t \mid X_{t-1}^{a_t}, \widetilde{\Xi}'_t). \tag{17}$$

Given the current values (a'_t, Ξ'_t), a Markov kernel with (13) as its stationary distribution can be constructed as in Algorithm 3. The validity of this approach follows as a special case of the derivation of the PG kernel [3].

Algorithm 3 Conditional importance sampling

Require: Current state (a'_t, Ξ'_t).

1: Set $\breve{a}_t^N = a'_t$ and $\breve{\Xi}_t^N = \Xi'_t$.

2: Draw $(\breve{a}_t^i, \breve{\Xi}_t^i)$ from (17) for $i = 1, \ldots, N-1$.

3: Set

$$\breve{w}_{t-1}^i = \frac{w_{t-1}^{\breve{a}_t^i}\gamma_T(X_{t-1}^{\breve{a}_t^i} \cup \breve{\Xi}_t^i \cup \widetilde{\Xi}'_t)}{\nu_{t-1}^{\breve{a}_t^i}\psi_t(\breve{\Xi}_t^i \mid X_{t-1}^{\breve{a}_t^i}, \widetilde{\Xi}'_t)\gamma_{t-1}(X_{t-1}^{\breve{a}_t^i})}$$

for $i = 1, \ldots, N$.

4: Draw ℓ with $\mathbb{P}(\ell = i) \propto \breve{w}_{t-1}^i$.

5: **return** $(a^\star_t, \Xi^\star_t) = (\breve{a}_t^\ell, \breve{\Xi}_t^\ell)$.

Remark 5. We can also define the kernel K_t to be composed of m, say, iterates of the MH or the CIS kernel to improve its mixing speed (at the cost of an m-fold increase in the computational cost of simulating from K_t). Indeed, any standard combination of MCMC kernels (see, e.g., [45]) targeting (13) will result in a valid definition of K_t.

4.4 Convergence properties

Existing convergence analysis for particle Gibbs algorithms [32, 4, 10] can be extended also to the proposed modified PGAS procedure of Algorithm 2. Here, we restate the uniform ergodicity result for the PGAS algorithm presented by [31], adopted to the current settings. We write $\|\cdot\|_\infty$ and D_{TV} for the supremum norm and the total variation distance, respectively.

Theorem 1. *Assume that there exists a constant $\kappa < \infty$ such that $\|\omega_t\|_\infty \leq \kappa$ for any $t \in \{1, \ldots, T\}$. Then, for any $N \geq 2$ there exist constants $R_N < \infty$ and $\rho_N \in [0, 1)$ such that*

$$D_{\mathrm{TV}}(\mathrm{Law}(X_T[k]), \bar{\gamma}_T) \leq R_N \rho_N^n, \qquad \forall X'_T \in \mathsf{X}^T,$$

where the Markov chain $\{X_T[k]\}_{k\geq 0}$ is generated by iterating Algorithm 2 with initial state $X_T[0] = X'_T$.

The proof follows analogously to the proof of [31, Theorem 3] and is omitted for brevity.

5 Degenerate and nearly degenerate models

The method proposed in Algorithm 2 can be used for a general sequence of target distributions $\{\gamma_t(X_t)\}_{t=1}^T$. We now turn our attention explicitly to (nearly) degenerate SSMs as described in **(M1)** and discuss how Algorithm 2 can be used for these models.

5.1 Ancestor sampling for nearly degenerate models

Consider again inference for the SSM given by (1), with the unnormalised target density $\gamma_t(X_t) = p(X_t, Y_t)$. We follow the convention used in Algorithms 1 and 2, that the reference particle is always placed on the Nth position. It follows that the ancestor sampling distribution (8) is given by

$$\mathbb{P}(a_t^N = i) = \frac{w_{t-1}^i f(x'_t \,|\, x_{t-1}^i)}{\sum_{j=1}^N w_{t-1}^j f(x'_t \,|\, x_{t-1}^j)}, \qquad i = 1, \ldots, N. \tag{18}$$

If the state process noise is small, i.e. the transition density $f(\cdot)$ is nearly degenerate, then this probability distribution can be highly concentrated on $i = N$, effectively removing the effect of ancestor sampling; we experienced this effect in the motivating example in Section 4.1.

To cope with this issue, one option is to make a partial collapse over a subset of the future state variables. That is, we let $\Xi_t = (x_t, \ldots, x_{\kappa_t})$ where $\kappa_t = \min\{T, t+\ell-1\}$ for some fixed length ℓ. It follows that the ratio of the unnormalised target densities appearing in (13) can be written as

$$\frac{\gamma_T(X_T)}{\gamma_{t-1}(X_{t-1})} = p(\widetilde{X}_t, \widetilde{Y}_t \,|\, x_{t-1}) \propto f(x_{\kappa_t+1} \,|\, x_{\kappa_t}) \left\{ \prod_{s=t}^{\kappa_t} f(x_s \,|\, x_{s-1}) g(y_s \,|\, x_s) \right\}. \tag{19}$$

Hence, the target distribution for the modified ancestor sampling step, with particle rejuvenation, is defined on $\{1, \ldots, N\} \times \mathsf{X}^{\kappa_t - t + 1}$ and the corresponding PDF of (a_t, Ξ_t) is proportional to

$$w_{t-1}^{a_t} f(x'_{\kappa_t+1} \,|\, x_{\kappa_t}) \left\{ \prod_{s=t+1}^{\kappa_t} f(x_s \,|\, x_{s-1}) g(y_s \,|\, x_s) \right\} f(x_t \,|\, x_{t-1}^{a_t}) g(y_t \,|\, x_t). \tag{20}$$

Simulating from this distribution can be done, e.g., by using one of the MCMC kernels introduced in Section 4.3. The benefit of doing this is that by rejuvenating the reference trajectory over the variables $(x_t, \ldots, x_{\kappa_t})$ we are able to bridge between x_{t-1}^i and x'_{κ_t+1}, thereby increasing the probability of changing the ancestry for the reference path.

We used this approach with $\ell = 1$ and a CIS Markov kernel for state rejuvenation for the target tracking model in Section 4.1. Below, we present two other example applications where the aforementioned technique can be useful.

5.2 Example: Euler-Maruyama discretisation of SDEs

Consider a continuous-time state space model with hidden state $\{Z(\tau)\}_{\tau \geq 0}$, represented by the following stochastic differential equation (SDE)

$$dZ_\tau = \mu(Z_\tau) d\tau + \sigma(Z_\tau) dW_\tau, \tag{21}$$

where W_τ denotes a Wiener process. The process is observed indirectly through the observations $(y_1, \ldots, y_T)$, obtained at time points $(\tau_1, \ldots, \tau_T)$, where $y_t \sim g(y_t \,|\, Z_{\tau_t})$. A simple approach to enable inference in this model is to consider a time discretisation of the continuous process, using an Euler-Maruyama scheme, after which standard discrete-time inference techniques can be used. For simplicity, assume that the observations are equidistant, $\Delta\tau := \tau_t - \tau_{t-1}$, and that we sample the process m times for each observation.

Let the discrete-time state at time τ_t consist of Z_{τ_t}, as well as the $m-1$ intermediate states, i.e.

$$x_t := \begin{pmatrix} \bar{x}_{t,1}^\mathsf{T} & \cdots & \bar{x}_{t,m-1}^\mathsf{T} & Z_{\tau_t}^\mathsf{T} \end{pmatrix}^\mathsf{T} \tag{22}$$

where $\bar{x}_{t,j} = Z_{\tau_{t-1}+j\Delta\tau}^\mathsf{T}$ for $j = 1, \ldots, m-1$. When using PGAS for this model, a problem is that, while increasing m makes the discretisation more accurate, it will also make the transition kernel of the latent process more degenerate.

However, this issue can be mitigated by rejuvenating the intermediate state variables, i.e. the state variables in between observation time points. Hence, we set

$$\Xi_t = \begin{pmatrix} \bar{x}_{t,1}^{\mathsf{T}} & \cdots & \bar{x}_{t,m-1}^{\mathsf{T}} \end{pmatrix}^{\mathsf{T}}. \tag{23}$$

Similarly to above, it follows that the ratio of unnormalised target densities is given by

$$\begin{aligned} \frac{\gamma_T(X_T)}{\gamma_{t-1}(X_{t-1})} &\propto f(x_t \,|\, x_{t-1}) && (24) \\ &\propto p(Z_{\tau_t} \,|\, \bar{x}_{t,m-1}) \left\{ \prod_{j=2}^{m-1} p(\bar{x}_{t,j} \,|\, \bar{x}_{t,j-1}) \right\} p(\bar{x}_{t,1} \,|\, Z_{\tau_{t-1}}), && (25) \end{aligned}$$

where, by the Euler-Maruyama discretisation,

$$p(Z_{\tau+\Delta\tau} \,|\, Z_\tau) \approx \mathcal{N}(Z_\tau + \mu(Z_\tau)\Delta\tau, \sigma^2(Z_\tau)\Delta\tau). \tag{26}$$

Hence, the unnormalised target PDF in (13) is given by

$$w_{t-1}^{a_t} p(Z'_{\tau_t} \,|\, \bar{x}_{t,m-1}) \left\{ \prod_{j=2}^{m-1} p(\bar{x}_{t,j} \,|\, \bar{x}_{t,j-1}) \right\} p(\bar{x}_{t,1} \,|\, Z_{\tau_{t-1}}^{a_t}). \tag{27}$$

To obtain an efficient MCMC proposal distribution for this PDF, we can use one of the methods proposed by [18] which are based on a tractable diffusion bridges between $Z_{\tau_{t-1}}^{a_t}$ and Z'_{τ_t}.

5.3 Example: Degenerate Gaussian transition

Consider a model with a Gaussian transition, but a possibly nonlinear/non-Gaussian observation

$$\begin{aligned} x_{t+1} &= Ax_t + Fv_{t+1}, && (28a) \\ y_t &\sim g(y_t \,|\, x_t), && (28b) \end{aligned}$$

with $v_t \sim \mathcal{N}(0, I_d)$. Furthermore, assume that $\dim(x_t) = n$ and that $\operatorname{rank}(F) < n$. This implies that the transition kernel of the linear Gaussian state process is degenerate. Models on this form are common in certain application areas, e.g., navigation and tracking; see [26] for several examples.

In this case, the ancestor sampling step is even more problematic than for the previous example, since the transition is truly degenerate. Indeed, if we use the ancestor sampling distribution from (8),[1] the probability of selecting x_{t-1}^i as the ancestor of x'_t will be zero, unless $x'_t - Ax_{t-1}^i$ is in the column space of F. However, in general, this will almost surely not be the case, except for $i = N$.

[1] Recall that we use density notation also for the degenerate kernel.

Thus, the distribution (8) puts all probability mass on $i = N$, resulting in zero probability of changing the ancestry of the reference trajectory. In fact, it is not only for PGAS the degeneracy of the transition kernel is problematic. As discussed in [30, Section 4.6], any conventional SMC-based forward-backward smoother will be inapplicable for the model (28) due to the degeneracy of the backward kernel.

However, by collapsing over intermediate state variables, this problem can be circumvented. We assume that the pair (A, F) in (28) are controllable (see e.g., [28] for a definition). Informally, this means that any state in the state space is reachable from any other state, i.e. for any $(x, x') \in \mathsf{X}^2$ there exists an integer ℓ and a noise realisation $v_{t:t+\ell}$ which takes the system from $x_{t-1} = x$ at time $t-1$ to $x_{t+\ell} = x'$ at time $t+\ell$. Now, let $\Xi_t = (x_t, \ldots, x_{\kappa_t})$ with $\kappa_t = \min\{T, t+\ell-1\}$ and where the length ℓ is chosen as any integer (e.g., the smallest) such that the matrix

$$C_\ell = \begin{bmatrix} F & AF & \cdots & A^\ell F \end{bmatrix}$$

is of rank n (the existence of such an integer is guaranteed by the controllability assumption).

Assuming $\kappa_t < T$ (the case $\kappa_t = T$ follows analogously) we have that the unnormalised target PDF in (13) is given by

$$w_{t-1}^{a_t} \left\{ \prod_{s=t}^{t+\ell-1} g(y_s \,|\, x_s) \right\} p(x'_{t+\ell}, \Xi_t \,|\, x_{t-1}^{a_t}), \tag{29}$$

where $p(x'_{t+\ell}, \Xi_t \,|\, x_{t-1}^{a_t})$ corresponds to the prior distribution under the linear Gaussian dynamics (28a) of the state sequence Ξ_t and the end-point $x'_{t+\ell}$, conditionally on the starting point $x_{t-1}^{a_t}$. Even though this distribution is degenerate for the model (28), our choice of ℓ ensures that for any $(x_{t-1}, x_{t+\ell}) \in \mathsf{X}^2$, there exists a $\Xi_t \in \mathsf{X}^{\kappa_t - t+1}$ in the support of the distribution.

In particular, the conditional distribution $p(\Xi_t \,|\, x_{t-1}^{a_t}, x'_{t+\ell})$ is a (degenerate) Gaussian distribution which it is possible to sample from. For instance, this can be done by running a Kalman filter/backward simulator [8, 22] for time steps $t, \ldots, t+\ell-1$ for the state process (28a), with $x'_{t+\ell} = Ax_{t+\ell-1} + Fv_{t+\ell}$ acting as an "observation" at the final time step (i.e., we condition on the reference state $x'_{t+\ell}$ via a standard measurement update of the Kalman filter). Consequently, it is possible to use this as a proposal distribution in the MH kernel (15) or in the CIS kernel (17) to simulate from (29). Indeed, in taking the ratio between the target (29) and the proposal $p(\Xi_t \,|\, x_{t-1}, x_{t+\ell})$ we have $p(x_{t+\ell}, \Xi_t \,|\, x_{t-1}) / p(\Xi_t \,|\, x_{t-1}, x_{t+\ell}) = p(x_{t+\ell} \,|\, x_{t-1})$ which is a well-defined Gaussian density for any $(x_{t-1}, x_{t+\ell}) \in \mathsf{X}^2$. Specifically,

$$p(x_{t+\ell} \,|\, x_{t-1}) = \mathcal{N}(x_{t+\ell} \,|\, A^{\ell+1} x_{t-1}, C_\ell C_\ell^{\mathsf{T}}), \tag{30}$$

which is non-degenerate under the controllability condition.

6 Near-degenerate approximations of intractable transitions

Another class of SSMs which poses large inferential challenges are models with intractable transition density functions as explained under label **(M2)**. Hence, consider an SSM on the form (1) and assume that the transition density function $f(\cdot)$ is a regular, non-degenerate PDF which it is possible to simulate from, but which is not available for evaluation in closed form. A problem with these models is that the backward kernel (2) is also intractable, essentially ruling out any forward-backward-based inference technique. In fact, one of the main merits of the PMCMC samplers derived in [3] is that, in their most basic implementations, they only require forward simulation of the system dynamics. These methods—specifically, PG and the particle independent Metropolis-Hastings (PIMH) sampler—can thus be readily used for inference in models with intractable transitions. However, these methods (PG and PIMH) are liable to poor mixing unless a large number of particles are used in the underlying SMC samplers (see, e.g., [31]). Intuitively, the reason for this is that we require the SMC sampler to generate approximate draws from the full *joint smoothing distribution*, which is difficult using only forward-simulation due to path space degeneracy.

As has been demonstrated (here and in the previous literature, e.g., [31, 30]), the PGAS sampler will in many cases enjoy much better mixing than PG and PIMH, in particular when using few particles N relative to the number of observations T. However, the PGAS sampler is not directly applicable to models with intractable transitions. Indeed, to simulate from the ancestor sampling distribution (18) it is necessary to evaluate the (intractable) transition PDF $f(\cdot)$. In this section, we propose one way to address this limitation. The idea is to approximate, as detailed below, the intractable transition PDF with a nearly degenerate transition. Using the proposed particle rejuvenation procedure, much in the same way as discussed in Section 5, we can then enable PGAS for this challenging class of SSMs.

The proposed method is essentially a variant of the approximate Bayesian computation (ABC) technique [5, 44]. However, while ABC is typically used for inference in models with intractable *likelihoods* (see, e.g., [11, 34] for SMC implementations), we use it here to address the issue of intractable transitions. The idea is based on the realisation that simulating from the transition PDF $f(\cdot)$, which is assumed to be feasible, is always done by generating some "driving noise variable" v_t, say, which is then propagated through some function $\Gamma(\cdot)$ (note that this function can be implicitly defined by a computer program or simulation-based software). By explicitly introducing these noise variables, we can thus rewrite the original model on the equivalent form,

$$v_t \sim p_v(v_t), \tag{31a}$$
$$x_t = \Gamma(x_{t-1}, v_t), \tag{31b}$$
$$y_t \sim g(y_t \mid x_t). \tag{31c}$$

Note that x_t here is given by a deterministic mapping of x_{t-1} and v_t. Consequently, the transition function for the *joint state* (v_t, x_t) is given by

$$p(v_t, x_t \mid v_{t-1}, x_{t-1}) = p_v(v_t)\delta_{\Gamma(x_{t-1}, v_t)}(x_t), \tag{32}$$

which is a degenerate transition kernel due to the Dirac measure on x_t.

Remark 6. The reformulation given by (31) can be seen as transforming the difficulty of having an intractable transition, to that of having a degenerate one. Now, if we marginalise over $x_{1:t}$ (which is straightforward since $x_{1:t}$ is deterministically given by $v_{1:t}$) we obtain a model specified only in the noise variables $\{v_t\}_{t\geq 1}$. This approach has previously been used by Murray et al. [36] in the context of auxiliary SMC sampling and for particle marginal Metropolis-Hastings. It was also used by Lindsten et al. [31] to enable inference by PGAS in a model with an intractable transition. The problem with that approach, however, is that the marginalisation of the x_t-process introduces a non-Markovian dependence in the observation likelihood, resulting in a T^2 computational complexity for PGAS (see [31] for details).

Simply rewriting the model as in (31) does not solve the problem, since, as discussed above, the degenerate transition kernel is problematic when using PGAS. However, by making use of an ABC approach this issue can be addressed. Specifically, we make use of a near-degenerate approximation of (32). In the ancestor sampling step of the algorithm we replace the point-mass distribution by some (for instance, Gaussian) kernel $\kappa_\epsilon : \mathsf{X}^2 \mapsto \mathbb{R}_+$ centered on $\Gamma(x_{t-1}, v_t)$:

$$\delta_{\Gamma(x_{t-1}, v_t)}(x_t) \approx \kappa_\epsilon(\Gamma(x_{t-1}, v_t), x_t), \tag{33}$$

where ϵ controls the band-width of the kernel (and thus the approximation error). Next, to deal with the near-degeneracy of the approximation (for small ϵ) we select $\Xi_t = v_t$ to be rejuvenated, which results in a joint target for (a_t, v_t), as in (13), proportional to

$$w_{t-1}^{a_t} p_v(v_t)\kappa_\epsilon(\Gamma(x_{t-1}^{a_t}, v_t), x_t'). \tag{34}$$

The connection to ABC is perhaps most easily seen if we make use of the CIS kernel given in Algorithm 3 for simulating from this distribution. Let $\psi_t(v_t \mid X_{t-1}, \widetilde{\Xi}_t') = p_v(v_t)$ be the proposal distribution for the noise variable (in many cases this is likely to be the only sensible choice) and let $\nu_{t-1}^{a_t} = w_{t-1}^{a_t}$ in (17). We then obtain the following ancestor sampling procedure, using the convention $a_t' = N$ (cf. Algorithm 3):

- For $i = 1, \ldots, N-1$:
  - Simulate $\breve{a}_t^i$ with $\mathbb{P}(\breve{a}_t^i = j) \propto w_{t-1}^j$.
  - Simulate a noise realisation $\breve{v}_t^i \sim p_v(\cdot)$ and set $\breve{x}_t^i = \Gamma(x_{t-1}^{\breve{a}_t^i}, \breve{v}_t^i)$.
  - Compute $\breve{w}_{t-1}^i = \kappa_\epsilon(\breve{x}_t^i, x_t')$.

- Compute $\breve{w}_{t-1}^N = \kappa_\epsilon(x'_t, x'_t)$.
- Simulate ℓ with $\mathbb{P}(\ell = i) \propto \breve{w}_{t-1}^i$, $i = 1, \ldots, N$.
- If $\ell < N$, return $a_t^\star = \breve{a}_t^\ell$, otherwise return $a_t^\star = N$.

In the above, we have assumed that the kernel approximation (33) is used only in the ancestor sampling step of the algorithm (i.e., in the forward simulation of particles we use the original model (31)). An interesting implication of this is that there is no need to explicitly keep track of the v_t-variables. Indeed, the CIS procedure outlined above can be expressed in words as follows: *(i)* Generate an independent set of $N-1$ resampled particles at time $t-1$ and, for each one, simulate the system dynamics forward to obtain $\{\breve{x}_t^i\}_{i=1}^{N-1}$, *(ii)* set the final particle according to the conditioning $\breve{x}_t^N = x'_t$, and *(iii)* simulate a new ancestor for x'_t based on the closeness of $\{\breve{x}_t^i\}_{i=1}^N$ to x'_t, as measured by the kernel κ_ϵ.

Remark 7. In some cases, for instance if X is high-dimensional, it can be beneficial to define the kernel κ_ϵ in (33) on some summary statistic $S : x_t \mapsto S(x_t)$, rather than on the state variable itself; see, e.g., [21] for details.

7 Numerical illustration

In this section we illustrate the particle rejuvenation strategy for PGAS on two examples. We have already seen the merits of the approach when compared to standard PGAS (and PG) on a nearly degenerate target tracking model in Section 4.1. Hence, here we consider two alternative model classes, first a model with a linear Gaussian degenerate transition as discussed in Section 5.3 and then a model with an intractable transition as discussed in Section 6.

7.1 Degenerate transition model

Autoregressive models are widely used to model stochastic processes [23]. They may be written in state space form as a degenerate Gaussian transition model,

$$x_{t+1} = \underbrace{\begin{bmatrix} \alpha_1 & \alpha_2 & \ldots & \alpha_{n-1} & \alpha_n \\ 1 & 0 & \ldots & 0 & 0 \\ 0 & 1 & \ldots & 0 & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & \ldots & 1 & 0 \end{bmatrix}}_{=A} x_t + \underbrace{\begin{bmatrix} \sigma_v \\ 0 \\ 0 \\ \vdots \\ 0 \end{bmatrix}}_{=F} v_{t+1}, \tag{35}$$

where $\alpha = \begin{pmatrix} \alpha_1 & \cdots & \alpha_n \end{pmatrix}^\mathsf{T}$ is a vector of regression parameters, $\{v_t\}_{t\geq 1}$ is (scalar) white Gaussian noise, and σ_v is the process noise standard deviation. In this example, we model the latent state as an autoregressive process of order $n = 5$. We simulate the system for $T = 500$ time steps using

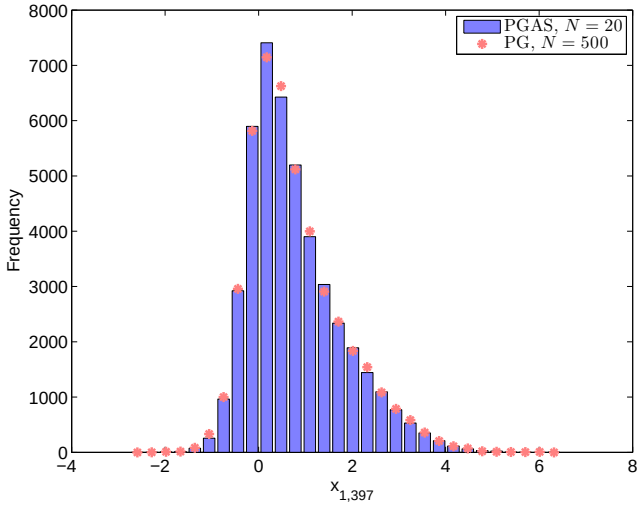

Figure 3: Posterior histograms for PGAS with $N = 20$ and $\ell = 4$ rejuvenated states (blue bars) and for PG with $N = 500$ (pink asterisks) for a randomly chosen state, $x_{1,397}$.

$\alpha = \begin{pmatrix} 0.9 & -0.8 & 0.7 & -0.6 & 0.5 \end{pmatrix}^{\mathsf{T}}$ and $\sigma_v = 1$. Each observation is a noisy, saturated measurement of the first component of x_t, modelled as,

$$y_t = \beta^{-1} \tanh(\beta x_{1,t}) + \sigma_e e_t, \tag{36}$$

where e_t is t-distributed with $\nu = 3$ degrees of freedom. We set $\beta = 0.5$ and $\sigma_e = 0.5$.

Since $\operatorname{rank}(FF^{\mathsf{T}}) = 1 < n$, the transition kernel is degenerate. Consequently, standard ancestor sampling is ineffective, and PGAS without rejuvenation is equivalent to basic PG. However, by collapsing over states $\Xi_t = (x_t, \ldots, x_{t+\ell-1})$ using $\ell \geq 4$, it is possible to update the particle ancestry. We use the CIS Markov kernel with new ancestor indexes sampled proportional to the filter weights, and state sequences sampled according to $p(\Xi_t \,|\, x_{t-1}, x_{t+\ell})$. The resulting CIS weights (see Algorithm 3, Step 3) are then,

$$\breve{w}_{t-1}^i = \left\{ \prod_{s=t}^{t+\ell-1} g(y_s \,|\, \breve{x}_s^i) \right\} p(x'_{t+\ell} \,|\, x_{t-1}^{\breve{a}_t^i}), \tag{37}$$

where $p(x_{t+\ell} \,|\, x_{t-1})$ is a well-defined Gaussian density given by (30).

We run the PGAS sampler with $\ell = 4$ rejuvenated states and with $N = 20$ particles. To check that the sampler indeed converges to the correct posterior distribution we also run a PG sampler with $N = 500$ particles (this value was chosen by trial-and-error as the smallest number required by PG to still have reasonable mixing). In Figure 3 we plot the histograms for the two samplers for

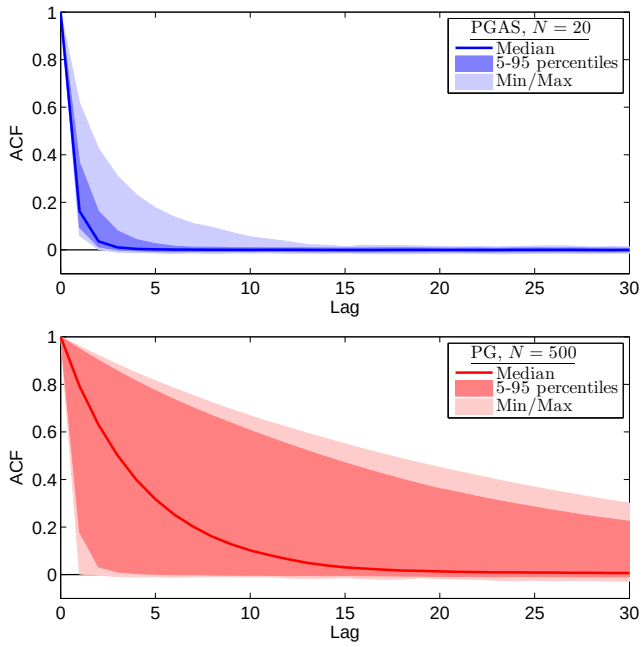

Figure 4: Empirical autocorrelations for the state variables $\{x_{1,t}\}_{t=1}^{500}$ for PGAS with $N = 20$ and $\ell = 4$ rejuvenated states (top) and for PG with $N = 500$ (bottom). The median (thick line), 5-95 percentile (shaded area), and min/max-values (lighter shaded area) over the 500 state variables are reported.

a randomly chosen state variable, $x_{1,397}$. As can be seen, there is a close match between the posterior histograms.

We also compute the empirical autocorrelation functions for both samplers for all state variables $\{x_{1,t}\}_{t=1}^{500}$. The results are reported in Figure 4. Despite the fact that it uses much fewer particles, the mixing speed of PGAS is significantly better than for PG. This is in agreement with previous results reported in the literature [31]. Indeed, the current example should mainly be seen as an illustration of how particle rejuvenation opens up for using backward-sampling-based methods, in particular PGAS, for a model where that would otherwise not be possible.

7.2 Intractable transition model

We now turn to a model with an intractable transition density to illustrate the ABC approximation for the PGAS sampler presented in Section 6. We consider inference in a stochastic version of the Lorenz '63 model [33], given by the

following SDE:

$$d\begin{bmatrix} Q_\tau \\ R_\tau \\ S_\tau \end{bmatrix} = \begin{bmatrix} \sigma(R_\tau - Q_\tau) \\ Q_\tau(\rho - S_\tau) - R_\tau \\ Q_\tau R_\tau - \beta S_\tau \end{bmatrix} dt + \begin{bmatrix} \sigma_Q & 0 & 0 \\ 0 & \sigma_R & 0 \\ 0 & 0 & \sigma_S \end{bmatrix} dW_\tau, \tag{38}$$

where W_τ is a three-dimensional Wiener process and the system parameters are $\sigma = 10$, $\rho = 28$, $\beta = 8/3$ and $\sigma_Q = \sigma_R = \sigma_S = \sqrt{5}$. The state is observed indirectly through noisy observations of the Q-component at regular time intervals: $y_t \sim \mathcal{N}(q_{t\Delta\tau}, 1)$ with $\Delta\tau = 0.01$. The initial state is distributed according to $(Q_0, R_0, S_0)^\mathsf{T} \sim \mathcal{N}(0, I_3)$.

A system simulator is implemented based on a fine-grid Milstein discretisation [35]. While the Milstein density for a single discretisation step is available [20], it is intractable to integrate out the intermediate steps on the grid. Consequently, the employed simulator lacks a closed form transition density function and, indeed, for the purpose of this illustration it is viewed simply as a "black-box" simulator.

We simulate the system for $\tau \in [0, 10]$ and thus generate $T = 1\,000$ observations $(y_1, \ldots, y_T)$. We then run PGAS with particle rejuvenation and the ABC approach outlined in Section 6 to compute the posterior distribution of the system state at the observation time points. The method uses $N = 100$ particles and a Gaussian kernel for the ABC approximation:

$$\kappa_\epsilon(x, x') = \exp\left(-\frac{\|x - x'\|^2}{2\epsilon}\right).$$

We let the kernel bandwidth range from $\epsilon = 0.01$ to $\epsilon = 10$. As comparison, we also run both the PG and PIMH samplers from [3] with the number of particles N ranging from 200 to 10 000 (the computational cost per iteration is roughly the same for PGAS with $N = 100$ as for PG/PIMH with $N = 200$, as the main computational cost comes from the system simulator).

RMSEs for the posterior means of the system states are shown in Figure 5.[2] The bias coming from the ABC approximation is evident for large ϵ, as the RMSEs level out at a non-zero value (if no approximation were made we would expect that the RMSE goes to zero[3] as the number of MCMC iterations increases). Nevertheless, comparing the results for PGAS to those obtained for PG and PIMH, it is evident that the ABC bias is significantly smaller than the Monte Carlo errors resulting from the poor mixing of PG and PIMH (at least if ϵ is not overly large).

As ϵ decreases the bias diminishes, but at the expense of slower convergence of PGAS. The reason for this is that the probability of updating the ancestry decreases with ϵ. In fact, for $\epsilon = 0$ the bias is completely removed, but we will then have zero probability of changing the ancestry and PGAS will be equivalent to PG (and thus suffer from the same poor convergence speed). Comparing the

[2] The "ground truth" is computed as an importance sampling estimator based on 10 000 independent particle filters, each one using $N = 50\,000$ particles.

[3] At least up to the accuracy of the "ground truth" reference sampler.

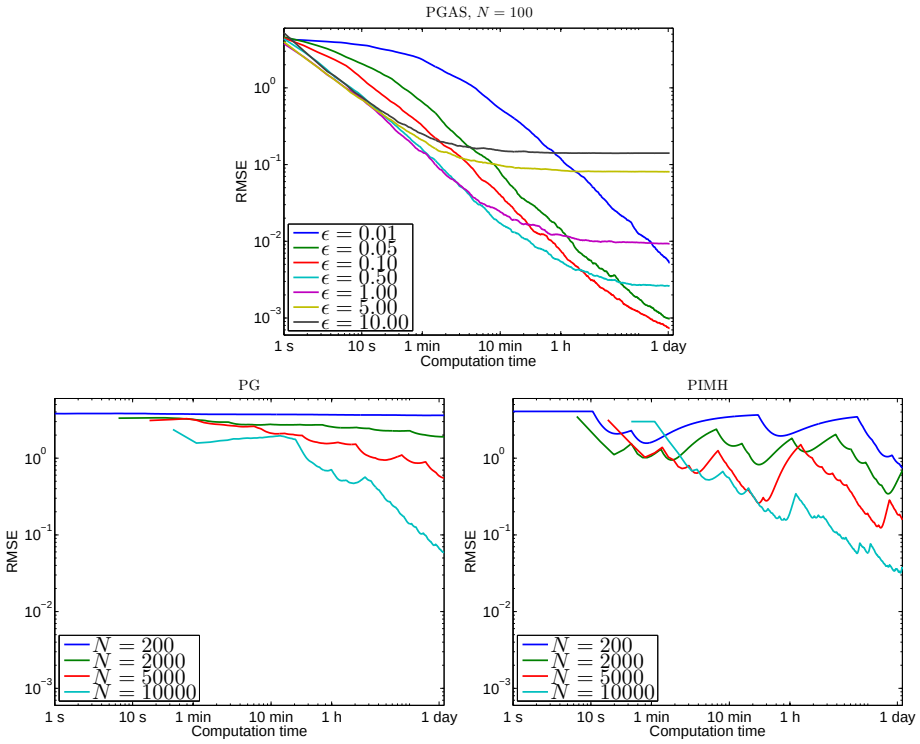

Figure 5: RMSEs in the estimated posterior mean $\mathbb{E}[x_{1:T} \mid y_{1:T}]$ for the Lorenz '63 model for PGAS using $N = 100$ (top), PG (bottom left), and PIMH (bottom right). (This figure is best viewed in color.)

results for PGAS using $\epsilon = 0.01$ with PG, however, we see that just having a small chance of updating the ancestry can have a significant impact on the mixing speed. For such a small value of ϵ the ABC bias is clearly dominated by the variance, even after 24 hours of simulation, corresponding to roughly 100 000 MCMC iterations.

8 Discussion

The particle rejuvenation technique presented in this paper generalises existing backward-simulation-based methods and opens up for a high degree of flexibility when implementing these procedures. This flexibility has been shown to be crucial for obtaining efficient samplers for several challenging types of state space models with (nearly) degenerate and/or intractable transitions. However, the technique is more generally applicable and we believe that it can be useful also for other types of models. In fact, we have recently made use of the particle rejuvenation technique in a completely different setting, namely to prove the validity of the nested SMC algorithm presented in [37]. To further investigate the scope and usefulness of the particle rejuvenation technique in other contexts is a topic for future work.

Our main focus in this paper has been on state smoothing (or, more generally, inference for a latent stochastic process). However, one of the main strengths of PMCMC samplers, such as the PGAS algorithm that we have used as the basis for the presented technique, is that they can be used for *joint state and parameter* inference. For PGAS, this is typically done by implementing a two-stage Gibbs sampler by iterating:

1. Simulate the parameter θ from its full conditional given the states X_T and observations Y_T.
2. Simulate the states X_T from the PGAS Markov kernel, conditionally on θ and Y_T.

This approach can be used also with the proposed Algorithm 2 to obtain a valid MCMC sampler for the joint posterior $p(\theta, X_T \mid Y_T)$. Indeed, this is the method that we used to sample the parameter θ in the illustrative example in Section 4.1. However, it is worth pointing out that this approach might not always be successful for the challenging model classes **(M1)** and **(M2)** that have largely motivated the present development. The problem is that for these models it *could* be infeasible to simulate θ from its full conditional in Step 1 of the aforementioned Gibbs sampler, since the degeneracy or intractability of the transition density *could* be inherited by full conditional distribution of θ. In such scenarios, we thus need a different way for enabling the use of Algorithm 2 for parameter inference. We mention here two possible, albeit as of yet untested, approaches.

Firstly, even if the model is degenerate or intractable, it is typically possible to explicitly introduce the "noise variables" $\{v_t\}_{t=1}^T$ that drives the state transition; see (31). We can then design a Gibbs (or Metropolis-within-Gibbs) sampler

for the extended model, with the original state variables X_T marginalised out. Note that we can still use Algorithm 2 to simulate X_T, but then transform the states to $\{v_t\}_{t=1}^T$ when updating θ. This overcomes the prohibitive $O(T^2)$ computational complexity associated with using PGAS for simulating the driving noise variables *directly*, as discussed in [31].

Secondly, it is possible to couple Algorithm 2 with the particle marginal Metropolis-Hastings (PMMH) algorithm by Andrieu et al. [3]. PMMH simulates (θ, X_T) jointly and implements a Metropolis-Hastings accept/reject step based on an estimate of the data likelihood computed by running a (forward-in-time only) SMC sampler. A problem with PMMH, however, is that the method tends to get "stuck", due to occasional overestimation of the likelihood. We believe that the method proposed in this paper can be used to mitigate this issue. Indeed, it is possible to use Algorithm 2 to refresh the likelihood estimate used in PMMH, while still maintaining the correct limiting distribution of the sampler (the details are omitted for brevity). By occasionally refreshing the likelihood in this way, it may thus be possible to escape the sticky states with overestimated likelihoods that deteriorate the practical performance of PMMH.

Investigating the effectiveness of these approaches, as well as enabling parameter inference in models of types **(M1)** and **(M2)** by using the method presented in Algorithm 2 in a more direct sense, are topics for future work.

Another interesting and important direction for future work is to analyse the effect of the ABC approximation (33). It was found empirically in [31] that PGAS appears to be robust to approximation errors in the ancestor sampling weights, and this is in agreements with our findings reported in Section 7.2. However, a more theoretical analysis is called for to understand if the sampler affected by the ABC approximation still admits a limiting distribution and, if so, how this distribution is affected by the approximation error.

A Partially collapsed Gibbs sampler

The original PGAS method, reviewed in Algorithm 1, corresponds to the following partially collapsed Gibbs sampler for the extended target distribution (10); see [31]: Given $X_T = \widetilde{X}_1 = X'_T \in \mathsf{X}^T$ and $B_T = \widetilde{B}_1 = (N, \ldots, N) \in \{1, \ldots, N\}^T$:

(i) Draw $\mathbf{x}_1^{-b_1} \sim \pi_T^N(\,\cdot \,|\, \widetilde{X}_1, \widetilde{B}_1)$,

(ii) For $t = 2$ to T, draw:

(a) $(\mathbf{x}_t^{-b_t}, \mathbf{a}_t^{-b_t}) \sim \pi_T^N(\,\cdot \,|\, \mathbf{X}_{t-1}, \mathbf{A}_{t-1}, \widetilde{X}_t, \widetilde{B}_{t-1})$,

(b) $a_t^{b_t} \sim \pi_T^N(\,\cdot \,|\, \mathbf{X}_{t-1}, \mathbf{A}_{t-1}, \widetilde{X}_t, \widetilde{B}_t)$,

(iii) Draw $k \sim \pi_T^N(\,\cdot \,|\, \mathbf{X}_T, \mathbf{A}_T)$.

Similarly, the proposed PGAS algorithm with particle rejuvenation presented in Algorithm 2 corresponds to the following partially collapsed Gibbs sampler for (10):

(i) Draw $\Xi_1 \sim \pi_T^N(\,\cdot\,|\,\widetilde{\Xi}_1, \widetilde{B}_1)$,

(ii) Draw $\mathbf{x}_1^{-b_1} \sim \pi_T^N(\,\cdot\,|\,\widetilde{X}_1, \widetilde{B}_1)$,

(iii) For $t = 2$ to T, draw:

(a) $(\mathbf{x}_t^{-b_t}, \mathbf{a}_t^{-b_t}) \sim \pi_T^N(\,\cdot\,|\,\mathbf{X}_{t-1}, \mathbf{A}_{t-1}, \widetilde{X}_t, \widetilde{B}_{t-1})$,

(b) $(a_t^{b_t}, \Xi_t) \sim \pi_T^N(\,\cdot\,|\,\mathbf{X}_{t-1}, \mathbf{A}_{t-1}, \widetilde{\Xi}_t, \widetilde{B}_t)$,

(iv) Draw $k \sim \pi_T^N(\,\cdot\,|\,\mathbf{X}_T, \mathbf{A}_T)$.

More precisely, Ξ_1 and $(a_t^{b_t}, \Xi_t)$ are sampled from the Markov kernels K_1 and K_t, respectively, in Steps *(i)* and *(iii-b)*. However, since these Markov kernels are constructed to leave the corresponding conditional distributions invariant, this corresponds to a standard composition of MCMC kernels. The fact that the Gibbs sampler outlined above is properly collapsed, and thus leaves π_T^N invariant, follows by analogous arguments as in the proof of [31, Theorem 1].

References

[1] C. Andrieu and M. Vihola. Convergence properties of pseudo-marginal Markov chain Monte Carlo algorithms. *The Annals of Applied Probability (forthcoming)*, 2012. arXiv:1210.1484.

[2] C. Andrieu, N. de Freitas, A. Doucet, and M. I. Jordan. An introduction to MCMC for machine learning. *Machine Learning*, 50(1):5–43, 2003.

[3] C. Andrieu, A. Doucet, and R. Holenstein. Particle Markov chain Monte Carlo methods. *Journal of the Royal Statistical Society: Series B*, 72(3):269–342, 2010.

[4] C. Andrieu, A. Lee, and M. Vihola. Uniform ergodicity of the iterated conditional SMC and geometric ergodicity of particle Gibbs samplers. arXiv.org, arXiv:1312.6432, December 2013.

[5] M. A. Beaumont, W. Zhang, and D. J. Balding. Approximate Bayesian computation in population genetics. *Genetics*, 162(4):2025–2035, 2002.

[6] P. Bunch and S. Godsill. Improved particle approximations to the joint smoothing distribution using Markov chain Monte Carlo. *IEEE Transactions on Signal Processing*, 61(4):956–963, 2013.

[7] P. Bunch, F. Lindsten, and S. S. Singh. Particle Gibbs with refreshed backward simulation. In *Proceedings of the 40th IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, Brisbane, Australia, 2015. (accepted for publication).

[8] C. K. Carter and R. Kohn. On Gibbs sampling for state space models. *Biometrika*, 81(3):541–553, 1994.

[9] C. K. Carter, E. F. Mendes, and R. Kohn. An extended space approach for particle Markov chain Monte Carlo methods. arXiv.org, arXiv:1406.5795, July 2014.

[10] N. Chopin and S. S. Singh. On particle Gibbs sampling. *Bernoulli*, 2014. Forthcoming.

[11] T. A. Dean, S. S. Singh, A. Jasra, and G. W. Peters. Parameter estimation for hidden Markov models with intractable likelihoods. *The Scandinavian Journal of Statistics*, 41(4):970–987, 2014.

[12] P. Del Moral. *Feynman-Kac Formulae - Genealogical and Interacting Particle Systems with Applications*. Probability and its Applications. Springer, 2004.

[13] P. Del Moral, A. Doucet, and A. Jasra. Sequential Monte Carlo samplers. *Journal of the Royal Statistical Society: Series B*, 68(3):411–436, 2006.

[14] P. Del Moral, R. Kohn, and F. Patras. On particle Gibbs Markov chain Monte Carlo models. arXiv.org, arXiv:1404.5733, 2014.

[15] A. Dempster, N. Laird, and D. Rubin. Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society, Series B*, 39(1): 1–38, 1977.

[16] A. Doucet and A. Johansen. A tutorial on particle filtering and smoothing: Fifteen years later. In D. Crisan and B. Rozovskii, editors, *The Oxford Handbook of Nonlinear Filtering*, pages 656–704. Oxford University Press, Oxford, UK, 2011.

[17] A. Doucet, M. K. Pitt, G. Deligiannidis, and R. Kohn. Efficient implementation of Markov chain Monte Carlo when using an unbiased likelihood estimator. *Biometrika (forthcoming)*, 2014. Preprint, arXiv:1210.1871v3.

[18] G. B. Durham and A. R. Gallant. Numerical techniques for maximum likelihood estimation of continuous-time diffusion processes. *Journal of Business & Economic Statistics*, 20(3):297–316, 2002.

[19] D. A. Van Dyk and T. Park. Partially collapsed Gibbs samplers: Theory and methods. *Journal of the American Statistical Association*, 103(482):790–796, 2008.

[20] O. Elerian. A note on the existence of a closed form conditional transition density for the Milstein scheme. Economics Discussion Paper 1998-W18, Nuffield College, Oxford, 1998.

[21] P. Fearnhead and D. Prangle. Constructing summary statistics for approximate Bayesian computation: semi-automatic approximate Bayesian computation. *Journal of the Royal Statistical Society: Series B*, 74(3):419–474, 2012.

[22] S. Frühwirth-Schnatter. Data augmentation and dynamic linear models. *Journal of Time Series Analysis*, 15(2):183–202, 1994.

[23] S. Godsill and P. Rayner. *Digital audio restoration*. Springer, 1998.

[24] S. J. Godsill, A. Doucet, and M. West. Monte Carlo smoothing for nonlinear time series. *Journal of the American Statistical Association*, 99(465):156–168, March 2004.

[25] A. Golightly and D. J. Wilkinson. Bayesian parameter inference for stochastic biochemical network models using particle Markov chain Monte Carlo. *Interface Focus*, 1(6):807–820, 2011.

[26] F. Gustafsson, F. Gunnarsson, N. Bergman, U. Forssell, J. Jansson, R. Karlsson, and P.-J. Nordlund. Particle filters for positioning, navigation, and tracking. *IEEE Transactions on Signal Processing*, 50(2):425–437, 2002.

[27] J. Handschin and D. Mayne. Monte Carlo techniques to estimate the conditional expectation in multi-stage non-linear filtering. *International Journal of Control*, 9(5):547–559, May 1969.

[28] T. Kailath, A. H. Sayed, and B. Hassibi. *Linear Estimation*. Prentice Hall, Upper Saddle River, NJ, USA, 2000.

[29] X R Li and V P Jilkov. Survey of maneuvering target tracking. part I: Dynamic models. *IEEE Transactions on Aerospace and Electronic Systems*, 39(4):1333–1364, 2003.

[30] F. Lindsten and T. B. Schön. Backward simulation methods for Monte Carlo statistical inference. *Foundations and Trends in Machine Learning*, 6(1):1–143, 2013.

[31] F. Lindsten, M. I. Jordan, and T. B. Schön. Particle Gibbs with ancestor sampling. *Journal of Machine Learning Research*, 15:2145–2184, 2014.

[32] F. Lindsten, R. Douc, and E. Moulines. Uniform ergodicity of the particle Gibbs sampler. *Scandinavian Journal of Statistics (forthcoming)*, 2015. doi: 10.1111/sjos.12136. Preprint, arXiv:1401.0683.

[33] E. N. Lorenz. Deterministic nonperiodic flow. *Journal of the Atmospheric Sciences*, 20(2):130–141, 1963.

[34] T. McKinley, A. R. Cook, and R. Deardon. Inference in epidemic models without likelihoods. *The International Journal of Biostatistics*, 5(1):1557–4679, 2009.

[35] G. N. Milstein. A method of second-order accuracy integration of stochastic differential equations. *Theory of Probability and its Applications*, 23:396–401, 1978.

[36] L. M. Murray, E. M. Jones, and J. Parslow. On disturbance state-space models and the particle marginal Metropolis-Hastings sampler. *SIAM/ASA Journal on Uncertainty Quantification*, 1(1):494–521, 2013.

[37] C. A. Naesseth, F. Lindsten, and T. B. Schön. Nested sequential Monte Carlo methods. arXiv.org, arXiv:1502.02536, February 2015.

[38] J. Olsson and T. Rydén. Metropolising forward particle filtering backward sampling and Rao-Blackwellisation of Metropolised particle smoothers. Technical Report 2010:15, Mathematical Sciences, Lund University, Lund, Sweden, 2010.

[39] M. K. Pitt, R. S. Silva, P. Giordani, and R. Kohn. On some properties of Markov chain Monte Carlo simulation methods based on the particle filter. *Journal of Econometrics*, 171:134–151, 2012.

[40] D. A. Rasmussen, O. Ratmann, and K. Koelle. Inference for nonlinear epidemiological models using genealogies and time series. *PLoS Comput Biology*, 7(8), 2011.

[41] C. P. Robert and G. Casella. *Monte Carlo Statistical Methods*. Springer, 2004.

[42] D. B. Rubin. A noniterative sampling/importance resampling alternative to the data augmentation algorithm for creating a few imputations when fractions of missing information are modest: The SIR algorithm. *Journal of the American Statistical Association*, 82(398):543–546, June 1987. Comment to Tanner and Wong: The Calculation of Posterior Distributions by Data Augmentation.

[43] M. A. Tanner and W. H. Wong. The calculation of posterior distributions by data augmentation. *Journal of the American Statistical Association*, 82(398):528–540, June 1987.

[44] S. Tavaré, D. J. Balding, R. C. Griffiths, and P. Donnelly. Inferring coalescence times from DNA sequence data. *Genetics*, 145(2):505–518, 1997.

[45] L. Tierney. Markov chains for exploring posterior distributions. *The Annals of Statistics*, 22(4):1701–1728, 1994.

[46] J. A. Vrugt, J. F. ter Braak, C. G. H. Diks, and G. Schoups. Hydrologic data assimilation using particle Markov chain Monte Carlo simulation: Theory, concepts and applications. *Advances in Water Resources*, 51:457–478, 2013.

[47] N. Whiteley. Discussion on Particle Markov chain Monte Carlo methods. *Journal of the Royal Statistical Society: Series B*, 72(3):306–307, 2010.

[48] N. Whiteley, C. Andrieu, and A. Doucet. Efficient Bayesian inference for switching state-space models using discrete particle Markov chain Monte Carlo methods. Technical Report Bristol Statistics Research Report 10:04, University of Bristol, 2010.